

## ***SAM (REX/Ops) Operations Meeting 2007.10.23***

Present: Adam (typist), Randolph, Art, SteveW, Anil, Robert, Krzysztof, SteveS, Andrew, Krzysztof, Dehong, Angela, Parag, Dave, ByeongRok (virtually)

For reports submitted to the meeting...

**The report text is in this font**

*Comments made during the meeting are in this font inline with the report*

### **Releases:**

sam\_bootstrap - v8\_1\_1 (SteveW) increases the lag time in the case of repeated db server restarts up to once an hour. Once the DB server has been running for more than two hours, the restart lag resets down to 3 seconds. Steve has installed it on all DØ nodes. This will reduce the huge number of log files problem.

No changes to address the problem of two weeks ago (when two DB servers started up upon a ups startup). Not sure what the problem is there. Can't easily reproduce it. Note that there is no protection against starting a new DB server when the old one is running when the PID file is changed or lost.

Typical problem is that someone does "ups stop sam\_bootstrap", that clears the PID file, but for some reason the DB server fails to stop. But this problem is somewhat rare and needs a design discussion and cost benefit analysis.

SteveW is working on new release for db\_server\_base\_cx. Robert wants to build the latest omniOrb (no ups kit yet). Steve's new code will allow the DB server to reregister with the nameservice (perhaps set to every hour). Due out soon.

### **DØ:**

**Kin Yip to sam-design**

**--- SamTV : D0 SAM shifters complain about its delay between updates often. Again, should we start to use SamHDTV ?**

**Then, why not put a link in the backdoor page ( <http://d0db.fnal.gov/sam/backdoor/> ) or so ?**

*Could switch over soon. SteveS will put in a project limit to keep HDTV from taking too long. Some little work to merge the web pages. Update the links to HDTV tomorrow morning. :-)*

**--- <http://www-d0en.fnal.gov/> (and below) was reported to be inaccessible on Oct. 16/17, 2007 .**

*Some web pages from d0en are always inaccessible from offsite (like the tape inventory). What we need is an off local network node that we have access to for checking these things. Perhaps when the new wireless subnets come up this will be possible easily. But note that there are some work arounds - see the sam shifter guide.*

*Note that the CD does offer VPN that makes your personal remote machine look like its onsite.*

No Unresolved issues left from the last week or almost last month.  
Woo hoo!

## **CDF:**

ByeongRok Ko to sam-design  
Dear all,

Here is cdf sam shift report

Cheers,  
cdf sam shifter

#####

Tue Oct 16 :

Many enstore problems

email from Ken Schumacher

There have been many LSM library failures on the SL-8500 at FCC. (really at GCC)  
The problem had been a tape stuck in the hand of Robot #1.  
That tape is one of many new blank tapes  
which were being inserted into robot.  
The Sun technician checked the hand.  
I used that Robot arm to unload tapes from four drives as a first  
test. Then the libraries were brought back on-line.

Wed Oct 17 :

pnfsmanager was offline

Monitoring issue,  
email from Stan

George tried the documented procedure. It does not work. We need a  
proper procedure. Also, how would we know this is broke if a customer  
had not called it in. What alarm should have been sent to notify  
admins that the pnfs manager is down.  
We can't continuously view the dcache monitor page.

IA7226 was showed up as a bad tape

email from George Szmuksta

There is problem with this tape. The robot has lost track of it.

==>They fixed it

Thu Oct 18 :

Tape IA7226 failed again.

email from George Szmuksta

The external label on the tape has been replaced.

==>It was fixed.

Mon Oct 22 :

r-fcdfdata320-1

r-fcdfdata320-2

r-fcdfdata320-3

r-fcdfdata320-4

were offline, but they were pingable.

email from Ling

/var is broken. Looks like we will have to reboot the machine.

Krzysztof: Machine has a problem with the OS installation and needs to be reinstalled.

*K: On the CAF: security scanning is affecting dcap client. When waiting for a callback, dcap client gets a response from the security scanning machine instead of from the dcache server. dcap client will sometimes recover, but will sometimes time out. Perhaps disable scanning? This is a new problem -- is the scanner now hitting this port? dcap client perhaps needs to handle this. Client called current in ups/upd also has this problem. There is a newer client, but not in ups yet (is available as rpm, but that doesn't work for cdf - must be ups'ified). Not known if this latest client will have the same problem. When the security scans are running a significant fraction of dcap clients fail. This is bad for the CAF.*

*Robert: Is gridftp safe from these scans? It also uses callbacks.*

*Art: Need a test from the security people to test the effects of scanning.*

*K: Are SAM clients and servers safe? We don't know yet. -- Andrew says that all communications are Corba, and Corba should handle this correctly (see comment below on corba)*

*Randolph: datacomm has told him that it is allowable for a service (or client) to blacklist an IP address that is sending malformed packets. So long as you don't*

*explicitly target the security scanners.*

*Robert: Corba and omniOrb seems to be robust against junk packets.*

*Dave: encp had a similar problem – was fixed by adding tags to the packets and having the client check.*

*K: CDF will need more tapes when the accelerator starts up.*

*Angela: There's a new version of fcp out.*

*Robert: Helpdesk ticket has been there for two weeks about a service principle with respect to SRM. Dave Berg: Cert should work through gplazma, but it does not. ???*

## **MINOS:**

Arthur Kreymer      to sam-design, minos\_sam\_admin

Minos status 2007 Oct 23

### General -

Deploying Condor on Minos Cluster interactive systems this week,  
startd's not registering with collector

### Enstore -

### DCache -

Improved root access to DCache, with 24 second startup overhead.

Using HEAD version of root

Using tree->SetCacheSize(50000000)

*Using the v5 versions of Root to access dCache is about 30 times slower than v4 Root. See above for work around (only with CVS head version of Root). There is overhead for the first file, but the MINOS framework has a 60s delay anyway when a file is opened. This workaround gives full speed access to dCache with Root!!*

### BlueArc -

Continuing to load datasets, no problems so far.

Got these mounted on FNALU batch, asked for GPARM also,  
consistent with the evolving OSG baseline security document.

*Not sure if these disks can be mounted on the open science enclave (and under what conditions). Awaiting plan from the CD.*

SAM -

IT 1894 Need 'sam locate' in C++ API

Oracle -

Scheduled minosora1 production Oracle upgrade to 10.2.0.3 and OS patches this Thursday 18 Oct.

*Upgrade went fine. No issues.*

## OPEN ISSUES

IT 1146: cannot test existence of sam tape location

IT 1894: sam locate via C++ api, for simple client needs

IT 1979: Compiler warnings from sam\_cpp\_api

IT 2071: quadratic project startup time growth

PLAN -

Migration to 9940B tape can proceed anytime - when ?

## DB:

Anil Kumar to sam-design

Oracle released the critical patch update on Oct 16.

We will be scheduling this patch on dev/int and prd databases.

D0 :

- Scheduling the Oracle's critical patch on d0ofdev1 on 10/23 Tuesday between 8:00AM-9:00AM *Already done - went well.*

- Scheduling the Oracle's critical patch on d0ofint1/d0ofprd1 on Nov 06 Tuesday

- SAM cut v6\_10 on d0ofdev1 - To add composite index on lower(file\_name) and retired\_date -done

- v6\_4 index on create\_date of data\_files against d0ofdev1 done

- Service Request with Oracle on Full table scan Vs Index scan on data\_files\_param\_values.

Development is working on Bug Fix. Bug# 5601578 No time estimate yet for bug fix.  
SAR# 5837876.993

Alerts:

**CDF :**

- Planning to refresh SAM schema @d0ofint1 on 10/24 Wed.

- Working with Randy on Query tuning.

**Alerts:**

Information for: cdfofpr2

Owner Object	Type	Extents Used	Extents Allocated	Extents Left
SAMPRD AP_PK	INDEX	96	100	4

Action : Increased the max extents to 256

**MINOS :**

- upgrade to 10.2.0.3 on minosprd and kernel patch - done on Oct 18.

Art doesn't want to take another downtime so soon to deploy the latest patch. Anil says that the patch can wait. Good.

*Randolph: Can we have a functional non-unique index on data\_files.file\_name (where the function is lower). This is to ensure that a new file being added doesn't have the same name as another file (with only case differences). Randolph thinks this index is necessary for efficiency (like thisIsAFile <--> THISISAFIle). Doing this check without the index requires a full table scan. Randolph also wants an insertion/update trigger to do the file name check (so that manual updates would be protected as well).*

*Lum\_min, lum\_max - the range'ness leads to inefficient queries. Perhaps make a many-to-many table. This is a sam design issue.*

**enstore/dCache:**

*The CDF tape mentioned above was stuck in a hand (that's the GCC robot, the FCC robot is not operational yet). Tape was removed and hand checked out ok.*

*System got in a state where tapes were in the "cap" cartridge access port" (bidirectional hole in the robot where humans can access tapes) and the robot was not removing them, for some reason. Sun engineer fixed.*

*9940B tape had a label in the wrong position. Robot intermittently read it correctly, and sometimes not (see CDF report above). Label was replaced.*

*SL8500 at FCC is waiting for a part that was damaged in shipping. Not an off-the-shelf item, so it may take awhile to arrive.*

*Fire protection system is being installed in the second robot at GCC (no drives in there yet). Perhaps plan is to install LTO4 drives in there. But since the two robots at GCC are connected, both robots must be at the same firmware. Therefore, the CMS tape labels must be upgraded (since only the latest firmware knows about LTO4s).*

*Plans for a downtime on this Thursday on STKN on the SL8500 to do the firmware upgrade, but that assumes the labeling was done, which it is not. So may cancel that downtime.*

*One pool node on public dCache is offline due to RAID issues. All precious files were drained. Can now investigate the machine.*

*Scheduled downtime for d0en on 11/6. Plan is to migrate the entire suite of d0en servers to a new rack of new servers and raids with SDE version of enstore (screwdriverless). Maintainability should be improved.*